Philos Stud (2011) 152:335–345
DOI 10.1007/s11098-009-9483-6

The A-theory of time and induction

Alexander R. Pruss

Published online: 24 November 2009
© Springer Science+Business Media B.V. 2009

Abstract The A-theory of time says that it is an objective, non-perspectival fact about the world that some events are *present*, while others *were* present or *will be* present. I shall argue that the A-theory has some implausible consequences for inductive reasoning. In particular, the presentist version of the A-theory, which holds that the difference between the present and the non-present consists in the present events being the only ones that exist, is very much in trouble.

Keywords Time · A-theory · Presentism · Induction

1 Introduction

The A-theory of time says that it is an objective, non-perspectival fact about the world that some events, like my writing this paper, are *present* (though by the time you read the paper, that event will no longer be present, but your reading it will be), while other events were present or will be present. I shall argue that the A-theory implausibly weakens inductive reasoning from past data to present or future conclusions, and thus should be rejected.

The A-theory comes in three major varieties. The moving spotlight (MSL) view allows that past, present and future events[1] and objects exist *simpliciter*, but says that there is a non-perspectival fact about which of the events and objects are *present*, and which ones are *past* or *future*. The growing block (GB) variety allows

[1] Unless, of course, the particular moving spotlight theorist has some independent reason to reject the existence of events altogether. One would have to hear what such a theorist's story about induction is before seeing how to adapt the criticism of this paper.

A. R. Pruss (✉)
Department of Philosophy, Baylor University, One Bear Place #97273, Waco, TX 76798-7273, USA
e-mail: alexander_pruss@baylor.edu

that past and present events and objects exist *simpliciter*, with there being a non-perspectival fact about which of them are past and which are present. But according to GB there are no future events and objects. Rather, reality is a growing block, constantly accreting new events and objects at its leading edge, namely the present. A GB-theorist may allow that there are truths about the future, but if these truths have truthmakers, the truthmakers exist in the present or past (maybe the truth of a proposition that some event will occur is grounded in the fact of an abstract state of affairs presently having the property of futurity).

Finally, the presentist is the most abstemious ontologically: there are only *present* events and objects. The typical presentist will, however, hold that there are truths about the past—Napoleon really was vanquished at Waterloo. These truths are either made true by facts about presently existing things, such as the abstract state of affairs of Napoleon being vanquished at Waterloo presently having the property *pastness*,[2] or have no truthmakers at all.[3] Some presentists will allow for truths about the future, either only in the case of those propositions whose truth value is determined by past and present events,[4] or to the extent of full bivalence.

The B-theorist, on the other hand, holds that while there are objective distinctions between the *earlier-than*, *later-than* and *simultaneous-with* relations, the question of what events or objects are *present* has only a perspectival answer.

My argument against the A-theory does not depend on scepticism about induction. Instead, it makes use of standard canons of inductive reasoning. I will formulate the argument first in terms of simple induction and then in terms of universal generalization.

2 Simple induction

I have observed many ravens. They were all black. I now find out that around the corner there is a raven named Claris. It seems I should assign a high probability to Claris being black. But suppose I learn that all the ravens which I have hitherto observed were male, while Claris is female. I should then at least somewhat decrease my confidence in Claris's being black. For I have learned of an additional difference between the instance I am making an inference about and the other instances I have observed.

Now, in every case of induction, there is at least one difference between observed (using the term widely) and unobserved instances, namely in that they are differently related to the observer. But that is a difference present in *every* piece of inductive reasoning, and hence has already been taken into account when I assigned my initial high probability to Claris being black. In fact, it may be that in induction mere differences in the relation between the cases and the *particular* inductive reasoner do not by itself count—indeed, that may be essential to the nature of inductive reasoning. The difference in sex between observed and unobserved

[2] For a discussions of approaches of this sort, see Bigelow (1996) and Crisp (2007).

[3] See Merricks (2007).

[4] E.g., Rhoda et al. (2006).

instances, however, is not like that—it affects some cases of induction but not others.

How much of a difference in probabilities the sex disparity makes depends on further inductive information. If I know that in most bird species males and females usually have *different* coloration, then the difference in sexes will decrease the probability of Claris being black to less than 1/2. If, on the other hand (and contrary to fact), I have good inductive evidence from other species that male and female birds always have the same coloration, then discovering the difference in sex will only slightly decrease the probability of Claris being black. However, even so it *will* decrease it, because now one is depending on *two* inductive arguments to establish Claris's blackness: first, the inductive argument about the blackness of ravens, and, second, the inductive argument about coloration not differing between the sexes.

In general, then, each time I learn of yet another difference between observed cases and the case I am interested in, the strength of the inductive argument goes down. Of course, the difference may enable *some other* inductive argument given additional evidence. Thus, if I had a very strong inductive argument that all female birds are black, then learning that Claris is a female raven might well increase the probability of Claris being black. It would weaken the inference to Claris's blackness from the blackness of the observed (male) ravens, but overcompensate for this by making possible an inference to Claris's blackness from the blackness of the observed (non-raven—for we have assumed no female ravens were observed) female birds.

Now, let us move to the A-theory case. Suppose I know about events $e_1,\ldots,e_n$ of type E, and also know that they all have feature F. I now learn of an additional event, h, of type E, which is later than all the e_i. I thus have an inductive argument in favor of h having F. How strong the argument will be depends, of course, on the details of the situation and data. But, and this is crucial, if I learn of any additional difference between h and the e_i, my induction will decrease in strength.

Suppose, then, that that I learn two facts: the first is that all the e_i are past, and the second is that h is present. (A variant of the argument is where I learn that h is future, and in response to the objections I will sometimes discuss the one variant and sometimes the other.) If the A-theory is true, this is an objective difference between h and the e_i (the B-theorist denies the objectivity of the difference). Therefore, I ought to lower my confidence in h's having F after learning that there is such a difference between h and the e_i. Thus, if the A-theory is correct, then the strength of a simple induction depends on how the cases are temporally positioned with respect to the present.

But this conclusion is false. For while it may matter in what temporal relations the events stand to each other (thus, if all the e_i are clustered close to each other temporally and h is far away from them, then the induction is weaker), and in what temporal relations the events stand to other events, it should not matter for inductive purposes whether the data for the induction is in the past and the case we are making an inference to is also in the past, or the data for the induction is in the past and the case we are interested in is in the present or in the future.

Since supposing the A-theory leads to a false view of induction—namely, the view that how cases are related to the present by itself affects the strength of an inductive argument—it follows that the A-theory is false.

Let us now consider some objections.

2.1 Objection 1

When we do inductive reasoning, in general the difference between (a) how the inductive data is related to the inductive reasoner and (b) how the case we are trying to draw a conclusion about is related to the inductive reasoner does matter. The results of some experiments after all are affected by whether they are observed, say for quantum mechanical reasons or because the subjects of the experiment notice that they are being observed. Therefore, the argument against the A-theory equally applies against the B-theory. For, h is differently temporally related to the inductive reasoner at the time of reasoning than $e_1,\ldots,e_n$ are, and hence the argument is not specifically an argument against the A-theory.

Response: But that was why I only suggested that differences in relation to the *particular* inductive reasoner do not matter for induction. Our inductive reasoning discounts the difference between cases observed by *us* and cases observed by *others*—that is a mere difference in the relation of the cases to the particular inductive reasoner. But inductive reasoning cannot always discount the difference between cases that are entirely unobserved and cases that are observed by someone or other. Indeed, if *this* difference were always discounted, the Copenhagen interpretation of quantum mechanics would be *clearly* untenable, the causal influences by observers (especially, but not only, in psychological experiments) would be neglected, and it might be difficult to account for selection effects such as that negative results are less likely to be reported.

Nonetheless, while the difference between an event's being entirely unobserved and its being observed by someone will lower the probability of the inductive conclusion, it will not defeat the inductive inference entirely. We were initially right to assume that *probably* micro-systems behave unobserved just as they behave when observed, *until* positive evidence to the contrary was found. However, even before the positive evidence came, we should have been more cautious about the inferences from observed cases to entirely unobserved ones than about inference from cases observed by us to cases observed by others.

2.2 Objection 2

Just as in drug testing trials, it is not usual to control for hair color because we know that hair color is an irrelevant feature, likewise we know that whether the inductive data is in the past, present or future does not in fact affect inductive arguments.

Response: However, three remarks need to be made in response. First, if a drug has worked flawlessly for thousands of patients and failed for no one, we will have very good reason to think that it will work for Genevieve. But if we learn that Genevieve is a brunette and none of the people on whom the drug was tested were brunettes, that would *somewhat* decrease our confidence that the drug would work

for her. After all, there really is a chance that there is some complex interaction between the drug and a correlate of hair color (perhaps explained by a common genetic cause, or by different social treatment of women of different colors of hair). We cannot rule out this possibility for certain, and hence we ought to slightly decrease our confidence in the drug working for Genevieve.

Second, our knowledge that hair color is irrelevant is itself inductive knowledge. Thus, while in the case of a someone who does not differ in hair color from all the patients on whom the drug was tested we need only one inductive argument—the argument from the data done on these people—in the case of a patient of different hair color, we need *both* this inductive argument and an inductive argument for the irrelevance of hair color to such tests. Thus, our confidence should be weaker in the case of Genevieve than in the case of someone who does not differ in hair color from all the tested patients.

Third, and perhaps most importantly, the parallel between hair color and A-properties like pastness, presentness and futurity, fails because A-properties, if there are any such, are ontologically important in ways in which hair color is not. It would be *prima facie* surprising if such an important ontological difference had no observational consequences. (That we think it has none is perhaps another argument against the A-theory.) The A-theory is an attempt to justice to a belief in the importance of the difference, a difference that A-theorists believe B-theorists do not take sufficiently seriously.

Perhaps something like an A-theory could be constructed on which the difference between past, present and future is entirely ontologically insignificant. Such an A-theory would not resemble any A-theory I know of. Here is one way to do this.[5] For every real number x, there is a certain basic property T_x that an event or an instant of time can have. All the properties T_x are metaphysically on par with one another, with none being metaphysically prior to or more natural than the others. Then, an event or instant is said to be "now" provided it has T_0, "future" provided there exists an $x > 0$ such that it has T_x, and "past" provided there exists an $x < 0$ such that it has T_x. Since there is nothing special about T_0—it's completely metaphysically on par with, say, $T_{120.74}$—the present is not ontologically different from the past. But, nonetheless, there is an objective fact about what events are now—they are the ones that have T_0. However, there is nothing special about the "0", though it may look special to us. To avoid that appearance, I could relabel the properties, letting T^*_x be T_{x-7}, and then an event would be present if and only if it had T^*_7. This qualifies as an A-theory by some definitions, but because it fails to do justice to the objective distinction between past, present and future, making it a merely numerical difference of no intrinsic significance, one would do better simply to opt for the B-theory. It is important to A-theories that the difference between past, present and future not be simply an arbitrary one.

[5] This is a modification of the referee's improvement on my original suggestion.

2.3 Objection 3

Maybe one's confidence in believing h has F should go down after learning that h is present while the e_i are past, but the difference in confidence is for all practical purposes insignificant. Our epistemic intuitions are unable to distinguish between small differences in credence, and so the intuition that A-properties are inductively irrelevant simply has to be rejected, with discomfort about the rejection assuaged by noting that the relevance is low. In other words, the A-theorist can bite the bullet, and the bullet is not so bad to bite.

Response: However, a judgment that a certain kind of difference matters little would have to be an a posteriori judgment standing in need of inductive justification. In our world, eye color is irrelevant to most bodily functions, so that, for instance, if we were testing a liver drug, we would not bother control for the patient's eye color. But this is an a posteriori judgment we make on the basis of inductive data—we have not observed a correlation between eye color and anything else of physiological significance when one controls for factors such as being an albino, and so we ignore eye color. We can easily imagine beings very much like us where eye color is much more significant.

Could one gather inductive evidence that differences in A-properties are not very relevant? Presumably, this would as follows. In the past, we have made many inductive inferences from the then-past to the then-present or then-future, and these inferences have all turned out to be just as good as inferences from the then-past to the then-past. Therefore, by induction, inferences from the past to the present or future should be just as good as inferences from the past to the past. I will call this reasoning the "A-meta induction".

One worry about the A-meta induction is that it is circular. We are inferring from past cases of engagement in inductive reasoning to present or future cases of engagement in inductive reasoning. However, I have not claimed that my argument *undercuts* reasoning from past to present or future cases given the A-theory, but only that it *weakens* such reasoning. Thus, for all that I have argued, the A-meta induction does give us reason to think that the difference in confidence between inferences from past to present or future and those from past to past is small.

Nonetheless, there will have to be a decrease in confidence. And any decrease in confidence here is absurd. Suppose I know that h is three hours in the future, and $e_1,\ldots,e_n$ are in the past. I am locked up in a dark room, far away from where h will be happening, and with no chance of ever observing h, and my only companion is a clock that I know to be accurate. I do the induction from the e_i's having F to h having F, and after noticing that h is in the future while the e_i are in the past, I slightly lower the probability I would otherwise have assigned to h's having F, and instead I assign p_0. I then wait, just watching the clock. Three hours later, h is present. Minutes later (depending on how long h takes), I know h is past. Now, the previously noted difference between h and the e_i disappears: they are all past events. So I increase the probability that h has F to some $p_1 > p_0$.

But isn't that absurd? With apparently no further relevant information gained, and without any of the tricks involved in Sleeping Beauty cases (such as forgetting, conditional waking up depending on what happened, etc.), I raise the probability

that *h* has *F*, exactly as I (we may suppose) expected I would have to. This is simply wrong.

Moreover, I think we have a strong *prima facie* presumption that A-properties are highly significant—at least A-theorists tend to think they are significant—and so the A-meta induction, because it itself depends on the claim that one can go from claims about the past to claims about the present or future, will not be that secure. If every first-order induction from past to present or future depends on the A-meta induction, then it will not give us any more confidence than the A-meta induction itself. But there are difficult issues here about whether induction can lend inductive support to itself, so perhaps it would be best to rest on the absurdity of raising the probability of *h* having *F* when the clock hits the right time.

This is perhaps all I can say to respond to Objection 3 in the case of moving spotlight views. But in the case of GB and especially presentism, a bit more can be said. According to GB, if I am making an inductive argument from past cases to a present one, I am making an argument to a case that stands literally *on the edge of time*, which is right at the boundary between the existent past and the non-existent future, from cases that do not have this property. We have good reason to think that events taking place at spatial boundaries are apt to differ significantly from those taking place away from the boundaries. If the earth were flat and had an edge, we would expect things to be different at the edge, and if the universe has a boundary, and if it were possible to be at the boundary, we would expect things to behave differently there. We do not at present know for sure if time has a beginning point, but if it does, our best candidate for that beginning point would be the Big Bang—and we do *not* expect things to behave at that point just as they do now. Boundaries *are* different.

And if I am making an argument from past cases to a *future* case, then GB will give an even more significant difference: I am making an argument from *existent* events $e_1,\ldots,e_n$ having property *F*, to a non-existent future event, *h*, having that property. It's not completely clear that this kind of talk of future events even makes sense, but perhaps the GB-theorist can say, along with some presentists, that propositions about the future are true in virtue of present states of affairs. In any case, something like this kind of talk of future events must make sense if GB is to have any hope of success (surely the GB-theorist has to be willing to say that there are *some* facts about the future). But then the difference between *existing* and *not existing* surely is *very* significant, much more so than the difference in hair color. Indeed, the difference may be so significant that we will judge it to be sufficient to defeat the inductive argument. (We would be rightly suspicious of an inductive argument from the claim that real detectives have some feature to the conclusion that Sherlock Holmes has this feature, though the nature of fictional characters may be too contentious for this analogy to bear much argumentative weight.)

And the presentist is particularly in a problematic situation when she is making an induction from what she takes to be *non-existent* past events to an *existent* present one. (An argument from the behavior of fictional detectives to the behavior of a real detective might be even more suspicious than an argument in the other direction.)

We might now offer the presentist or GB-theorist a dilemma. Given presentism or GB (respectively), either induction from past to present or future (respectively) is almost as good as induction from past to past, or it's not. If not, then the presentism or GB is in trouble. Suppose it *is* almost as good. Given the deep ontological difference between present and past on presentism, presentism gives us good *prima facie* reason to think that induction from past to present would be significantly worse than from past to past, and likewise GB gives us good *prima facie* reason to think that induction from past to future would be significantly worse than from past to past. But in fact induction from past to future is not significantly worse than induction from past to past. This is a *surprising* fact given presentism and GB. But it is not a surprising fact given the B-theory—it is exactly what we would expect given the B-theory. Thus, we have incremental confirmation of the B-theory against presentism and GB.

Finally, here is a curious little argument against GB. I do not know how good it is, but it is worth stating. Let us grant that induction from past to future is not much worse than induction from past to past. Then, it seems that from the fact that there are real sunrises in the past we can infer that there are real sunrises in the future. But if there are *real* future sunrises, then GB is false.

2.4 Objection 4

We have reason to be more confident about past things than future things. After all, there is always a possibility of an apocalyptic event, such as a global nuclear war, that could change many of our expectations. If h is in the past, the mere fact that I am alive gives me evidence to think that no such major changes have happened by h's time. But if h is in the future, in my inductive reasoning I may need to worry about apocalyptic scenarios. Even if h is present, I may need to worry about them, because perhaps the apocalyptic event has not yet reached me.

Response: This objection may force a small modification of the above cases. For instance, I might have to add to my background knowledge that I know for certain that I will survive past h without noticing any such global changes. How could I know that? Maybe God could give me a phone call and tell me, while working so many wonders around me (making a proof of the Riemann Zeta Conjecture materialize in mid air in glowing letters, regrowing a severed limb instantaneously, giving me a clearly sound ontological argument for his own existence, etc.) that I could be sure it was God. Or I might suppose that the dark room described in my response to Objection 3 is a place to which I have been supernaturally transported, a place beyond the scope of apocalyptic scenarios involving Earth, and h is an event on Earth, so my survival up to h's time is irrelevant.

But a better answer is this. We all have to take into account the possibility of an apocalyptic scenario in doing inductive reasoning. But the A-theorist has to give a *further* discount to the probabilities, on top of the apocalypse-discount, when h is in the present and future and all the e_i are in the past. That there is a further such discount is implausible.

3 Universal generalization

Again, go back to the case of the ravens. All the ravens in my sample are black. I form the universal generalization that all ravens are black. But I then learn that all ravens in the sample are male. I should now decrease the likelihood I assigned to the claim that all ravens are black. One way to explain this decrease is that when we know that all the ravens in the sample are male, we may opt for the more cautious hypothesis *all male ravens are black* instead of the stronger *all ravens are black*. If E is the event of all the ravens in the sample being black, then $1 = P(E \mid \text{all ravens are black}) > P(E \mid \text{all male ravens are black})$, and the hypothesis that all ravens are black receives more support from E than the hypothesis that all male ravens are black. But if E^* is the event of all the ravens in the sample being black and male, then $P(E^* \mid \text{all ravens are black}) = P(E^* \mid \text{all male ravens are black})$, at least to a high degree of approximation (assuming the sampling process is independent of the color and sex of the ravens), and so E^* does not give more support to the stronger hypothesis.

The same issue comes up in the A-theory case. All the events of type T in my sample have had property F. I form the hypothesis that all Ts are Fs. But when I learn that all the Ts in my sample were also past, if pastness is an objective property, I need to decrease the probability of the hypothesis that all Ts are Fs, as my data equally supports and would apparently be equally explained by the weaker claim that all past Ts are Fs.

Of course, when doing inductive reasoning, we need to avoid gerrymandered hypotheses, for instance because they have low prior probability. But the hypothesis that all male ravens are black is not gerrymandered, since maleness is quite a natural property of ravens. Similarly, while a hypothesis that all Ts before 4:36 pm on December 10, 2008 are Fs would be a gerrymandered hypothesis, because *being before 4:36 pm on December 10, 2008* is not at all a natural property of an event, it seems that according to the A-theory, *being past* is a natural property of an event given an A-theory (this is true even though at 4:36 pm on December 10, 2008, the two properties have the same extension).[6] Because of this, the B-theorist does not face the same problem as the A-theorist does.

In light of the above, the A-theory forces us to assign lower probabilities to a temporally universal generalization when we have inductive data all in the past than when we have inductive data in the past and the present. Moreover, the A-theory lowers our confidence in the temporally universal generalizations below the confidence we would have in them if the B-theory held even when we have inductive data in the past and in the present, because *non-future* is a fairly natural property given the A-theory (it's the negation of a natural property), and hence the hypothesis that all non-future Ts are Fs is pretty plausible. In other words, the A-theory weakens inductive reasoning from the past to the present and future in an implausible way.

Moreover, a problem of implausible change in probabilities arises here, similar to the discussion in the response to Objection 3. Let us suppose that I have observed all

[6] This will not be the case for the weird A-theory discussed at the end of the response to Objection 2.

*T*s being *F*s, and that my observations include some past *T*s and at least one present *T*. I form the temporally universal hypothesis that all *T*s are *F*s. But then a bit of time goes by, and the one present *T* that I had observed is now past, and I am not currently observing a present *T*. Now I should *decrease* my confidence in the temporally universal hypothesis. For while previously I had data from both past and present *T*s, now I only have data from past *T*s. Previously, I could not have supposed that only past *T*s are *F*s, but now I can. And this decrease in confidence is completely absurd. Hence, the A-theory should be rejected.

4 Conclusions

If the A-theory holds, then A-properties are relevant to inductive reasoning in ways in which they clearly are not relevant. Hence, the A-theory fails. The problem is greatest for GB and presentism, since there, the A-properties correspond to such radical ontological differences that it would be *surprising* if things in the future (in the case of GB) or the present (in the case of presentism) behaved as they did in the past.

Let me add one more note on why the problem is particularly great for presentism, in a rather sketchy way. Probably the best way to think of presentism[7] is to think that claims about the past and the future are like modal claims. The truth *simpliciter* of a proposition is its present truth. But just as a possibility operator applied to the proposition *that there are unicorns* yields the proposition that *possibly there are unicorns*, so there are WAS and WILL operators that applied to the same proposition yield the claims *that it was the case that there are unicorns* and *that it will be the case that there are unicorns*.[8] However, we now have a puzzle. Why should data within the scope of WAS and WILL operators be significantly inductively relevant to conclusions that lack any such operator? After all, *possibility* claims are of little inductive relevance: that every member of a sample of ravens is possibly black may give us good reason to think Claris is possibly black, but very little reason to think Claris is actually black. We acquire good reason to reject the claim that Claris simply *can't* be black, and that's some progress—there are entities, such as photons, that can't be black, and Claris isn't one of those, but we are still very far away from the claim that Claris actually *is* black. Likewise, why should accepting the proposition WAS(that all the ravens in the sample are black) give us much reason to think that Claris is black? The difference between claims within the scope of a possibility operator and claims outside it is, for all we know, immense. And so why should there not be a similar difference between claims in the scope of WAS and those outside of that scope?[9]

[7] Cf. Merricks (2007).

[8] Some would prefer to apply the operators to sentences, and the same point could still be made, but as a sociological matter of fact, presentists tend to be friendly to propositions.

[9] I would like to thank Jonathan Kvanvig and Trenton Merricks for helpful discussions of the argument of this paper, and Trent Dougherty for comments on an earlier version of this paper. I am very grateful to the referees for a careful reading and a number of objections that led to significant improvements in the paper.

References

Bigelow, J. (1996). Presentism and properties. *Philosophical Perspectives, 10*, 35–52.
Crisp, T. M. (2007). Presentism and the grounding objection. *Nous, 41*, 90–109.
Merricks, T. (2007). *Truth and ontology*. Oxford: Oxford University Press.
Rhoda, A. R., Gregory, G. A., & Belt, T. G. (2006). Open theism, omniscience, and the nature of the future. *Faith and Philosophy, 23*, 432–459.